

Raising a Hardness Result

Paolo Liberatore*

February 1, 2008 — 21:12

Abstract

This article presents a technique for proving problems hard for classes of the polynomial hierarchy or for PSPACE. The rationale of this technique is that some problem restrictions are able to simulate existential or universal quantifiers. If this is the case, reductions from Quantified Boolean Formulae (QBF) to these restrictions can be transformed into reductions from QBFs having one more quantifier in the front. This means that a proof of hardness of a problem at level n in the polynomial hierarchy can be split into n separate proofs, which may be simpler than a proof directly showing a reduction from a class of QBFs to the considered problem.

Contents

1	Introduction	2
2	Logic-Based Abduction	5
3	Default Logic	6
4	Planning	8
5	Conclusions	10

*Dipartimento di Informatica e Sistemistica Antonio Ruberti, Università di Roma “La Sapienza”, Via Ariosto 25, 00185, Roma, Italy. Email: paolo@liberatore.org

1 Introduction

Several logics-related problems are complete for classes of the polynomial hierarchy other than NP and coNP. This is because these problem involve logical consistency and entailment, which are already NP-hard and coNP-hard, respectively. Therefore, problems that are defined in terms of a number of consistency or entailment verifications may be complete for classes such as $\Delta_2^p[\log n]$, Δ_2^p , or even for higher classes when this number is exponential. An example of a problem that can be formulated in terms of an exponential number of such operations is that of checking the existence of a formula that is equivalent to a given one and of size bounded by a given number k . This problem can indeed be expressed as follows: if k is larger than the given formula, the answer is "yes"; otherwise, guess a formula of size bounded by k and check equivalence (mutual entailment) to the given formula. This problem is therefore in Σ_2^p . Hardness to the same class has been proved in particular cases [HW97].

Several results of Σ_2^p -hardness and Π_2^p -hardness have been published since the beginning of the 90s [Got92, Sti92, EG92, EG93, CL94, CDS94]. Some problems have even been proved to be hard for classes at the fourth level of the polynomial hierarchy [EG96, EGL97, DT02, ETW05].

For example, the problem of relevance in skeptical default abduction based on minimal explanations is Σ_4^p -complete. This has been proved by Eiter, Gottlob, and Leone [EGL97] by showing a reduction such that a QBF in the form $\exists X \forall Y \exists Z \forall K. F$ is valid if and only if a certain hypothesis is in some minimal explanations of a certain problem of skeptical default abduction. Such a proof is relatively complicated, as it requires showing that, if there exists an evaluation of the variables X such that for all possible evaluation of the variables Y , etc. then the hypothesis is in some minimal explanation, and vice versa. As a comparison, a proof of NP-hardness done by reduction from propositional satisfiability only involves an evaluation of the variables in X .

The technique proposed in this article simplifies such proofs by requiring only one quantifier at time to be considered. This may not necessarily simplify the search for a reduction, but allows its formal proof to be split into a number of simpler sub-proofs.

In particular, the idea is to start from the assumption that a reduction from QBF to a given problem works, and show that this reduction can be "raised", that is, modify in such a way the QBF has a single more quantifier in the front. For example, assume that a problem has been proved Π_3^p -hard by a reduction from QBFs in the form $\forall Y \exists Z \forall K. F$ to a given problem is already known. In some cases, we can use this reduction to produce a new reduction from QBFs in the form $\exists x \forall Y \exists Z \forall K. E$ to the problem under consideration, where x is a new variable. If this step can be iterated for a polynomial number of times, that would result in a proof of Σ_4^p -hardness. The (iterated) addition of a single quantifier

raised a Π_3^p -hardness proof to a Σ_4^p -hardness proof.

Formally, let P be problem under consideration, and assume that I is a translation from a class of QBFs into P . This means that every formula $Q.E$, where Q is a sequence of quantifiers and E a propositional formula, is translated into an instance $I(Q, E)$ of P such that:

$$Q.E \text{ is valid} \quad \text{iff} \quad I(Q, E) \in P$$

For any formula F containing a variable x , we let $F|_{x=v}$, where x is a variable and v is either **true** or **false**, be the formula obtained by replacing each occurrence of x with v . By this replacement, $F|_{x=v}$ does not contain the variable x . For example, $((y \wedge \neg x) \vee x)|_{x=\text{true}}$ is $(y \wedge \neg \text{true}) \vee \text{true}$. This formula is equivalent to **true** but syntactically different to it.

If F is a propositional formula containing only variables in Q and another variable x which is not in Q , then both $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$ are well-formed QBF formulae, since both $F|_{x=\text{true}}$ and $F|_{x=\text{false}}$ only contains variables in Q . Note that neither $Q.F|_{x=\text{true}}$ nor $Q.F|_{x=\text{false}}$ contain the variable x . Since $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$ have the same sequence of quantifiers Q , they can both be translated to P :

$$\begin{aligned} Q.F|_{x=\text{true}} \text{ is valid} & \quad \text{iff} \quad I(Q, F|_{x=\text{true}}) \in P \\ Q.F|_{x=\text{false}} \text{ is valid} & \quad \text{iff} \quad I(Q, F|_{x=\text{false}}) \in P \end{aligned}$$

Depending on the reduction I , the two instances $I(Q, F|_{x=\text{true}})$ and $I(Q, F|_{x=\text{false}})$ may be similar. If this is the case, one can try to merge them into a single instance $I'(Q, F, x)$ such that:

$$I'(Q, F, x) \in P \quad \text{iff} \quad \begin{cases} I(Q, F|_{x=\text{true}}) \in P \\ \text{or} \\ I(Q, F|_{x=\text{false}}) \in P \end{cases} \quad (1)$$

If such a merge is possible, it produces an instance $I'(Q, F, x)$ which is in P if and only if $\exists x Q.F$ is valid (note that F contains x):

$$\exists x Q.F \quad \text{iff} \quad I'(Q, F, x) \in P$$

If this step can be iterated a linear number of times while not super-polynomially increasing the size of the generated problem instances, then one has a reduction from the validity of $\exists X Q.F$ to P .

The key to the proof is Equation 1: the two instances of P can be merged. These two instances $I(Q, F|_{x=\text{true}})$ and $I(Q, F|_{x=\text{false}})$ are not arbitrary instances but the result of translating two QBFs with the same quantifier and similar matrixes.

This is often possible. In practice, many translations from QBFs to logic-based problems use the matrix of the QBF “as is”, by simply copying it verbatim in some part of the instance of the problem P . In this case, $I(Q, F|_{x=\text{true}})$ and $I(Q, F|_{x=\text{false}})$ are the same except for the part containing the matrix, where they only differ because one contains $F|_{x=\text{true}}$ and the other contains $F|_{x=\text{false}}$. In some cases, these two instances can be merged by simply taking $I(Q, F|_{x=\text{true}})$, replacing $F|_{x=\text{true}}$ with F , and minimally modifying the rest of the instance in such a way the answer can be expressed in terms of the answers to the two subproblems obtained by setting $x = \text{true}$ and $x = \text{false}$.

This merge needs not only to be possible, but also to generate instances such that merging can be applied again. If this iteration is possible while keeping the instance size polynomial, hardness can be raised of one level in the polynomial hierarchy. Consider a problem that has been proved Π_n^p -hard by a reduction from QBF to it. By iterating the step of adding an existential quantifier, one obtains a proof of Σ_{n+1}^p -hardness for the same problem.

This whole process may appear complicated at first, but is actually easier to perform to specific problems than to explain in its general form. One thing that one may easily overlook when considering specific problems is that of assuming that $Q.F$ is a QBF. This is not the case, as F contains the variable x , which is not in Q ; as a result, $Q.F$ is not a well-formed QBF. The QBFs mentioned in the proof are $Q.F|_{x=\text{true}}$, $Q.F|_{x=\text{false}}$, and $\exists x Q.F$. This is also reflected in the problem instances: $I(Q, F|_{x=\text{true}})$ and $I(Q, F|_{x=\text{false}})$ do not contain the variable x because x is not mentioned in Q , $F|_{x=\text{true}}$, and $F|_{x=\text{false}}$ (in the latter two formulae x is replaced by **true** and **false**, respectively.) The variable x only occurs in the instance $I'(Q, F, x)$.

A similar method can be used to prove that a universal quantifier can be added in front of a QBF. The only part that changes is that the merged instance is in P if both the two original instances are in P . In other words, the “or” in the right-hand side of Equation 1 is replaced by “and”.

In the following sections, we apply this technique to logic-based abduction [Pei55, BATJ91, EG95, CMP96, EGL97, EM02, LS07], default logic [Rei80, Got92, CS93, Ant99, BG02], and planning [FN71, BN92, Byl94, BJ95, Koe96]. In the first case, we add existential quantifiers, in the second universal quantifiers, and in the third both kinds.

2 Logic-Based Abduction

We consider the problem of checking the existence of explanations in logic-based abduction. This is in essence the problem of making hypotheses over the possible causes of observable manifestations [Pei55]. Formally, an instance of the problem of logic-based abduction is a triple $\langle H, M, T \rangle$, where H is a set of propositional variables (hypothesis), M another set of propositional variables (manifestations), and T a propositional formula relating H and M . An explanation is a subset $S \subseteq H$ such that $S \cup T$ is consistent and $S \cup T \models M$. Checking whether an explanation for an instance $\langle H, M, T \rangle$ exists is Σ_2^P -complete [EG95]. We provide an alternative proof using raising.

The starting point is a proof of coNP-hardness, which is easy to give: a formula E is inconsistent if and only if the following problem has explanations: hypotheses $H = \emptyset$, manifestations $M = \{a\}$, theory $T = \{\neg E \vee a\}$.

In this reduction, the formula is copied as is in the theory of the abduction problem. As a result, two formulae $E|_{x=\text{true}}$ and $E|_{x=\text{false}}$ are translated into two abduction instances differing only for the value of x of the theory T . We show that two such instances can merged into a single instance with a moderate increase of size.

Lemma 1 *For every H, M, T , the instance $\langle H', M', T' \rangle$ has exactly all explanations of $\langle H, M, T|_{x=\text{true}} \rangle$ with x^+ added to each and all explanations of $\langle H, M, T|_{x=\text{false}} \rangle$ with x^- added to each, where variable x does not occur in H and M , variables x^+, x^- , and q do not occur in H, M , and T , and:*

$$\begin{aligned} H' &= H \cup \{x^+, x^-\} \\ M' &= M \cup \{q\} \\ T' &= T \cup \{x^+ \rightarrow q, x^- \rightarrow q, x^+ \rightarrow x, x^- \rightarrow \neg x, \neg x^+ \vee \neg x^-\} \end{aligned}$$

Proof. Let S be an explanation of $\langle H', M', T' \rangle$. Since $q \in M'$, and this variable only occurs in the clauses $x^+ \rightarrow q$ and $x^- \rightarrow q$, then either $S \cup T \models x^+$ or $S \cup T' \models x^-$. Since x^+ and x^- does not occur positively in T' , this means that either $x^+ \in S$ or $x^- \in S$, but not both, since otherwise S would not be consistent with T' .

The explanations containing x^+ are exactly the explanations of $\langle H, M, T|_{x=\text{true}} \rangle$ with x^+ added to each. Indeed, $x^+ \in S$ makes $S \cup T'$ equivalent to $S \cup T|_{x=\text{true}} \cup \{\neg x^-, q, x\}$, and $T|_{x=\text{true}}$, H , and M do not contain x^- , q , and x . Similarly, the explanations containing x^- are exactly the explanations of $\langle H, M, T|_{x=\text{false}} \rangle$ with x^- added to each. As a result, the set of explanations of $\langle H', M', T' \rangle$ is (apart from x^- and x^+) the union of the explanations of $\langle H, M, T|_{x=\text{true}} \rangle$ and of $\langle H, M, T|_{x=\text{false}} \rangle$. \square

This lemma proves that two similar instances of abduction can be combined into a single one having the union of their explanations (apart from some variables added to each). As a result, if one is able to translate two QBFs $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$, then one can combine the resulting two instances into a single one that has explanations if and only if $\exists x Q.F$. In order to prove the hardness of the problem, one only needs to analyze the increase of size due to merging.

Theorem 1 (Alternative proof; originally proved by Eiter and Gottlob [EG95])
The problem of explanation existence is Σ_2^p -hard.

Proof. A QBF of the form $\forall Y.F$ is valid if and only if the problem of abduction $\langle \emptyset, \{a\}, \{\neg E \vee a\}$ has explanations. This reduction has the property that the matrix of the QBF is copied verbatim in the theory of the abduction instance.

Let us now assume the existence of a reduction with the same property from from QBFs having Q as their sequence of quantifier to abduction instances exists. If F is a formula made of variables of Q plus x , one can apply the previous lemma: $\langle H', M', T' \rangle$ has explanations if and only if either $\langle H, M, T|_{x=\text{true}} \rangle$ or $\langle H, M, T|_{x=\text{false}} \rangle$ has, where $\langle H, M, T|_{x=\text{true}} \rangle$ and $\langle H, M, T|_{x=\text{false}} \rangle$ are the results of translating $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$, respectively. As a result, $\exists x Q.F$ is valid if and only if $\langle H', M', T' \rangle$ has explanations.

What remains to be proved is that iterating this process does not generate instance of super-polynomial size. This is in this case straightforward, as each merge only adds a constant number of variables and binary clauses to the instance. \square

In this proof, it may look like $\langle H, M, T \rangle$ is an abduction instance involved in the proof. However, it is not. The instances used in the proof are $\langle H', M', T' \rangle$, $\langle H, M, T|_{x=\text{true}} \rangle$, and $\langle H, M, T|_{x=\text{false}} \rangle$. In other words, H and M are meant to be used only with $T|_{x=\text{true}}$ and $T|_{x=\text{false}}$ while T is meant to be used (with some formulae added to it) only with H' and M' .

More generally, the two instances corresponding to $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$ only contain the formulae $F|_{x=\text{true}}$ and $F|_{x=\text{false}}$, respectively, and not the formula F , which is instead contained in the merged instance. In particular, F contains the variable x , which is not mentioned in the two QBFs $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$; x is the variable of the merged instance that makes the merged instance become equivalent to one of the two original ones when assuming the value **true** or **false**.

3 Default Logic

In this section, we show how a reduction from QBF to the problem of skeptical entailment in default logic can be raised by the addition of a universal quantifier. This allows for

an alternative proof of Π_2^p -hardness of this problem. Default logic has been introduced by Reiter [Rei80]; several variants have been proposed since then [Luk88, Sch92, Ant99, Lib07]. The problem of checking whether a default theory skeptically entails a formula is Π_2^p -complete [Got92, Sti92].

The Π_2^p -hardness of a problem can be established by showing a reduction from $\forall\exists$ QBFs to the problem. The starting point is a simpler reduction from \exists QBF; this reduction is then raised by the addition of universal quantifiers. The starting reduction is easy to give: a propositional formula E is satisfiable if and only if $\langle \{ \frac{a \wedge E}{a \wedge E} \}, \emptyset \rangle$ skeptically entails a , where a is a variable not contained in E .

This is a reduction from \exists QBF to the problem of skeptical entailment in default logic. The QBF is translated in such a way its matrix only occurs once, as is, in the resulting defaults. Two default theories obtained by translating two QBFs having the same quantifiers and differing only for the value of a variable in the matrix can be merged as shown in the following lemma.

Lemma 2 *For every set of defaults D and variable q not occurring in D , the extensions of the following theory are exactly the extensions of $\langle D|_{x=\text{true}}, \emptyset \rangle$ with x and p added to each and the extensions of $\langle D|_{x=\text{false}}, \emptyset \rangle$ with $\neg x$ and p added to each.*

$$T = \left\langle \left\{ \frac{:xp}{xp}, \frac{:\neg xp}{\neg xp} \right\} \cup \left\{ \frac{p \wedge \alpha : \beta}{\gamma} \middle| \frac{\alpha : \beta}{\gamma} \in D \right\}, \emptyset \right\rangle$$

Proof. Since the first two defaults of T are mutually inconsistent, and they are the only ones that are applicable to the background theory, each extension of this default theory contains either $\{x, p\}$ or $\{\neg x, p\}$. The extension of T are therefore exactly the extensions of $\langle D, \{x, p\} \rangle$ and of $\langle D, \{\neg x, p\} \rangle$. In turn, these two theories have the same extensions of $\langle D|_{x=\text{true}}, \emptyset \rangle$ and of $\langle D|_{x=\text{false}}, \emptyset \rangle$, apart from x and q . \square

This lemma proves that two similar default theories can be merged into a single one having the extensions of both. Since skeptical entailment is considered, the latter theory implies a formula a if and only if both the two former theories do. If the two theories result from translating $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$, the merged theory therefore entails a if and only if $\exists x Q.F$ is valid.

Theorem 2 (Alternative proof; originally proved by Gottlob [Got92] and Stillman [Sti92])
Skeptical entailment in default logic is Π_2^p -hard.

Proof. A formula $\exists Y.E$ is valid if and only if $\langle \{ \frac{a \wedge E}{a \wedge E} \}, \emptyset \rangle$ skeptically entails a . In this reduction, the matrix of the QBF is copied verbatim in the justification of one default, and does not otherwise affect the default theory.

Let us now assume that a similar translation from QBFs having Q as their sequence of quantifiers to skeptical default entailment exists. We show a translation from QBFs having $\exists xQ$ as their sequence of quantifiers.

Let $\exists xQ.F$ be such a formula. By assumption, $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$ can be translated into two default theories where $F|_{x=\text{true}}$ and $F|_{x=\text{false}}$ only occur as the justification of a single default. As a result, the two theories corresponding to $Q.F|_{x=\text{true}}$ and $Q.F|_{x=\text{false}}$ can be written as $\langle D|_{x=\text{true}}, \emptyset \rangle$ and $\langle D|_{x=\text{false}}, \emptyset \rangle$, respectively, for some set of defaults D . One can then apply the previous lemma, which proves that these two theories both skeptically entail a if and only if T skeptically entails a . In other words, $T \models a$ if and only if $\exists xQ.F$ is valid. By construction, the translation from $\exists xQ.F$ to T has the same property that the matrix F is translated verbatim in a default.

In order to complete the proof, we calculate the increase of size of the involved default theories when the addition of quantifiers is iterated. This increase of size is that two new defaults (of constant size) are introduced, and a variable is added to the precondition of each default. As a result, this step can be iterated with only a quadratic increase of size, thus obtaining a reduction from $\exists\forall\text{QBF}$ to skeptical default entailment. \square

4 Planning

The problem of establishing the existence of a plan in STRIPS [FN71] is PSPACE-complete [BN92, Byl94]. We can use the method of raising for proving the hardness of this problem. In this case, we have to show that both $\exists x$ and $\forall x$ can be added to the front of a QBF while only producing a constant increase of size in the corresponding planning instance.

For the sake of simplicity, we consider an extension in which the precondition of each action is a propositional formula, rather than a list of positive and negative literals. An action is therefore a pair $\langle P, C \rangle$ where P is a formula and C is a list of literals. This action is executable if P is valid in the current state; its effect is to make all literals of C valid.

Clearly, checking whether a formula containing no variable is valid can be translated into a problem of plan existence. Given such a formula E , just build the action $\langle E, \{a\} \rangle$, where a is a (new) variable, and have a being false in the initial state and required to be true in the goal. This instance has a plan (composed of a single occurrence of the only action it contains) if and only if E evaluates to **true**.

This is a translation from QBFs with no quantifiers to the problem of planning. The matrix of the QBF is translated verbatim in the precondition of one action; the goal is a single variable. Let us now assume that such a translation from QBFs having Q as

their sequence of quantifiers exists, and prove the existence of a similar translation from QBFs with a single more quantifier in the front.

In order to add an existential quantifier $\exists x$, we add a new variable p which is initially false, we add two actions $\langle \neg p, \{x, p\} \rangle$ and $\langle \neg p, \{\neg x, p\} \rangle$, and add p as a precondition of all other actions. This way, p is required to be true before executing all other actions, which is only possible if x is made either true or false; once x has been given a value, p is also made true. This makes x unmodifiable, because the first two actions can no longer be executed, and no other action makes p false or changes the value of x . This way, a plan exists if and only if the instance corresponding to either $Q.F|_{x=\text{true}}$ or $Q.F|_{x=\text{false}}$ has a plan. The increase of size is of two new actions, plus one more literal in each action.

A universal quantifier is added as follows. Assume that x is the variable we want to quantify, and that the goal a is the postcondition of a single action. We add two new variables b and p , both false in the initial state, and change the goal from a to b . We also add p as a precondition to all other actions, and the following three actions.

$$\begin{aligned} a_1 &= \langle \neg p, \{x, p\} \rangle \\ a_2 &= \langle a \wedge x, \{\neg x, \neg a\} \rangle \\ a_3 &= \langle a \wedge \neg x, \{b\} \rangle \end{aligned}$$

In the initial state, only the first action is executable. It makes x true and all other actions executable. Let P be an irredundant plan of this instance. The last action of P is a_3 , since is the only action that makes the goal b true. This action requires both a to be true and x to be false. This means that, at some point, a_2 has been executed as well, since this is the only action that makes x false. As a result, we have that P starts with a_1 , contains a_2 , and ends with a_3 . Let P_1 and P_2 be the segments of P between a_1 and a_2 and between a_2 and a_3 , respectively. We have that P_1 is a sequence of actions that makes a true while x is true; P_2 makes a true while $\neg x$ is true. In other words, P_1 and P_2 are plans for the case in which x is assumed true and false, respectively.

All considered reductions have the following two properties: a. the matrix of the QBF only appears in the precondition of a single action and does not affect the rest of the planning instance; and b. the goal is a single variable. Given a reduction with these two properties, one can create a new reduction from QBFs with a single more quantifier in the front.

We omit the formal proof of PSPACE-hardness. We have informally proved that both an existential and a universal quantifier can be added to a QBF so that the corresponding planning instance only gain moderately in size. In particular, both changes add only two or three actions of constant size, and only add a single precondition to all other actions.

This means that the instance that results from adding n quantifier has at most $3 * n$ actions, and each action is at most $m + n$ large, where m is the size of matrix.

5 Conclusions

The method shown in this article allows for proving a result of hardness for a class of the polynomial hierarchy without directly building a reduction from the problem of validity of a class of QBFs. The idea is that a proof of Σ_n^p -hardness can be built by first showing that the problem is Π_{n-1}^p -hard, and then showing that the particular instances used in this hardness proof are able to “simulate” an existential quantifier. In the same way, one can prove Π_n^p -hardness from a Σ_{n-1}^p -hardness proof. This technique is shown here for explanation existence in logic-based abduction (Σ_2^p -hard), for skeptical default logic entailment (Π_2^p -hard), and plan existence in an extension of STRIPS (PSPACE-hard); in a previous article, it has been used to prove the Π_4^p -hardness of a problem related to redundancy in default logic [Lib05].

This technique can be used in two ways: either for all quantifiers or for the last one. For example, in the case of planning we have shown that planning instance can always simulate existential and universal quantifiers, therefore raising a Σ_0^p -hardness result to PSPACE-hardness. On the other hand, the proof of Σ_4^p -hardness of the problem of background theory redundancy in default logic [Lib05] is based on first showing the problem Π_3^p -hard using a “classical” reduction from QBFs, and then proving that existential quantifiers can be added, therefore raising this result to Σ_4^p -hardness.

The main advantage of this technique is the simplification of the proofs. While a proof of Σ_4^p -hardness would require considering a QBF in the form $\exists X \forall Y \exists Z \forall W. F$, raising an hardness result only requires to prove that two similar instances of the problem can be merged in the appropriate way. This means that considering complicated QBFs may not be necessary.

References

- [Ant99] G. Antoniou. A tutorial on default logics. *ACM Computing Surveys*, 31(4):337–359, 1999.
- [BATJ91] T. Bylander, D. Allemang, M. C. Tanner, and J. R. Josephson. The computational complexity of abduction. *Artificial Intelligence*, 49:25–60, 1991.

- [BG02] R. Baumgartner and G. Gottlob. Propositional default logics made easier: computational complexity of model checking. *Theoretical Computer Science*, 289(1):591–627, 2002.
- [BJ95] C. Bäckström and P. Jonsson. Planning with abstraction hierarchies can be exponentially less efficient. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI’95)*, pages 1599–1605, 1995.
- [BN92] C. Bäckström and B. Nebel. On the computational complexity of planning and story understanding. In *Proceedings of the Tenth European Conference on Artificial Intelligence (ECAI’92)*, pages 349–353, 1992.
- [Byl94] T. Bylander. The computational complexity of propositional STRIPS planning. *Artificial Intelligence*, 69:165–204, 1994.
- [CDS94] M. Cadoli, F. M. Donini, and M. Schaerf. Is intractability of non-monotonic reasoning a real drawback? In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI’94)*, pages 946–951, 1994. Extended version as RAP.09.95 DIS, Univ. of Roma “La Sapienza”, July 1995.
- [CL94] M. Cadoli and M. Lenzerini. The complexity of propositional closed world reasoning and circumscription. *Journal of Computer and System Sciences*, 48:255–310, 1994.
- [CMP96] M. Cialdea Mayer and F. Pirri. Abduction is not deduction-in-reverse. *Journal of the Interest Group in Pure and Applied Logic*, 4(1):86–104, 1996.
- [CS93] M. Cadoli and M. Schaerf. A survey of complexity results for non-monotonic logics. *Journal of Logic Programming*, 17:127–160, 1993.
- [DT02] B. Dimopoulos, Y. Nebel and F. Toni. On the computational complexity of assumption-based argumentation for default reasoning. *Artificial Intelligence*, 141:57–78, 2002.
- [EG92] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates and counterfactuals. *Artificial Intelligence*, 57:227–270, 1992.
- [EG93] T. Eiter and G. Gottlob. Propositional circumscription and extended closed world reasoning are Π_2^p -complete. *Theoretical Computer Science*, 114:231–245, 1993.

- [EG95] T. Eiter and G. Gottlob. The complexity of logic-based abduction. *Journal of the ACM*, 42(1):3–42, 1995.
- [EG96] T. Eiter and G. Gottlob. The complexity of nested counterfactuals and iterated knowledge base revisions. *Journal of Computer and System Sciences*, 53(3):497–512, 1996.
- [EGL97] T. Eiter, G. Gottlob, and N. Leone. Semantics and complexity of abduction from default theories. *Artificial Intelligence*, 90:177–223, 1997.
- [EM02] T. Eiter and K. Makino. On computing all abductive explanations. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI 2002)*, pages 62–67, 2002.
- [ETW05] T. Eiter, H. Tompits, and S. Woltran. On solution correspondences in answer-set programming. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI 2005)*, pages 97–102, 2005.
- [FN71] R. Fikes and N. Nilsson. STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
- [Got92] G. Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2:397–425, 1992.
- [HW97] E. Hemaspaandra and G. Wechsung. The minimization problem for Boolean formulas. In *Proceedings of the Thirtyeighth Annual Symposium on the Foundations of Computer Science (FOCS’97)*, pages 575–584, 1997.
- [Koe96] J. Koehler. Planning from second principles. *Artificial Intelligence*, 87:145–186, 1996.
- [Lib05] P. Liberatore. Redundancy in logic III: Non-monotonic reasoning. Technical Report cs.LO/0507048, Computing Research Repository (CoRR), 2005.
- [Lib07] P. Liberatore. Consistency defaults. *Studia Logica*, 86(1):89–110, 2007.
- [LS07] P. Liberatore and M. Schaerf. Compilability of propositional abduction. *ACM Transactions on Computational Logic*, 8, 2007.
- [Luk88] W. Lukasiewicz. Considerations on default logic: an alternative approach. *Computational Intelligence*, 4(1):1–16, 1988.

- [Pei55] C. S. Peirce. Abduction and induction. In J. Buchler, editor, *Philosophical Writings of Peirce*, chapter 11, pages 150–156. Dover, New York, 1955.
- [Rei80] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [Sch92] T. Schaub. On constrained default theories. In *Proceedings of the Tenth European Conference on Artificial Intelligence (ECAI'92)*, pages 304–308, 1992.
- [Sti92] J. Stillman. The complexity of propositional default logics. Presented at the 4th International Workshop on Nonmonotonic Reasoning, Vermont, May 1992. Notes edited by Kautz H. A. and Etherington D., 1992.